INTERNATIONAL NETWORK SERVICES

Constructing Scalable Frame Relay Networks with OSPF

Prepared by:
Tom Kunath
Senior Network Systems Consultant

INS Engineering Services
2000 Aerial Center, Suite 101
Morrisville, NC, 27560

Phone: (919) 319-0400
FAX: (919) 467-3650

Table of Contents

INTRODUCTION **3**

FRAME RELAY BACKGROUND **3**

FRAME RELAY COMPONENTS **4**

- The Network Access Circuit
- The Frame Relay Port
- The Virtual Circuit

FRAME RELAY PVC TOPOLOGY **6**

- PVC Meshing on a Network
- Partially Meshed Topology
- Fully Meshed Topologies

OSPF BACKGROUND **10**

OSPF BASICS **11**

- OSPF Neighbors
- OSPF Network Types
- OSPF Areas

DESIGN METHODOLOGY **15**

- Requirements Gathering
- Network Design
- Router Frame Relay PVC Modeling
- OSPF Design Guidelines for Frame relay Architectures
- OSPF Network Types

SUMMARY **22**

REFERENCES **23**

Introduction

The purpose of this white paper is to discuss strategies for designing and implementing networks using Open Shortest Path First (OSPF) as a TCP/IP routing protocol on frame relay wide area networks (WANs). The intent is to present an overview of the various technologies, as well as a design and implementation methodology with which efficient, scalable WANs can be built. Topics such as IP addressing, OSPF area organization, frame relay PVC architecture and router configuration guidelines will be the primary focuses of the discussion.

Frame Relay Background

Frame relay has rapidly become the data technology of choice for most corporate network WANs requiring speeds of T-1 (1.544MB) and below. Much of frame relay's popularity can be attributed to its deployment flexibility, efficiency as a data link protocol, public carrier availability, standards-based deployment, and economy of scale. Many public carriers now offer frame relay service speeds in increments of N x T-1, through multiplexing of multiple T-1 circuits. Recent modifications of frame relay standards have extended access speeds up to 45 MBPS (DS3), suggesting that this data link WAN protocol will be a survivor in a rapidly changing telecommunications world.

Efficiencies of this protocol are achieved through upper layer network and transport protocols for integrity checking, flow control, and network layer routing of data. Frame relay's speed and flexibility are maximized by removing any layer 3 functionality in the protocol. Networks built with frame relay utilize upper layer protocols such as TCP/IP and OSPF to find the appropriate and efficient path through the internetwork. While frame relay switches have mechanisms for detecting and adapting to periods of network congestion through forward and backward explicit congestion notifications, (FECN/BECN), most routers instead depend on end station host "sliding window" transport protocols such as TCP or SPX to detect congestion and throttle data rates accordingly.

Frame relay networks are considered to be non-broadcast multi-access (NBMA); i.e., frame relay switches will not duplicate broadcasts or multicasts from routers to PVCs in a virtual network. The burden of this duplication lies on the router: special challenges arise depending on the protocols carried, and routing protocols used to discover network layer addressing.

Frame Relay Components

Fundamental building blocks of frame relay networks include three major components: the network access circuit, the port, and the virtual circuit. Network equipment required to "ride" on frame relay networks includes frame relay access devices (FRADs), bridges, and multiprotocol routers. A brief discussion of the key components and equipment follows.

The Network Access Circuit:

The network access circuit is the data pipe that connects the customer network to the frame relay point of presence (POP). Supported methods include synchronous technologies such as N x DS0 (fractional T-1), DS1 (T-1), N x DS1 (multiplexed T-1), DS3 (45MB), and ISDN BRI or PRI. Data circuits are terminated on traditional DCE devices such as DSU/CSUs or ISDN terminal adapters connected to network equipment referenced above.

The Frame Relay Port:

The frame relay port is the interface point where the local loop meets the frame relay network. It can be literally mapped to a synchronous interface module and port on a particular frame relay switch. The frame relay port represents the maximum speed which data can enter (ingress point) or leave (egress point) a frame relay network. This maximum speed is often referred to as the maximum information rate (MIR) in carrier service level agreements. This is the maximum rate that data can burst to a virtual circuit.

The Virtual Circuit:

Virtual circuits represent the logical interconnections between frame relay switches within a network to provide end to end connectivity to specific local loop circuits and their associated network devices. Many virtual circuits (theoretically up to 1023) can be provisioned from a frame relay ingress port to various nodes within the network. Each virtual circuit is assigned a committed information rate (CIR), which represents the sustained rate at which traffic can flow end-to-end without being subject to discard during periods of network congestion.

Permanent virtual circuits (PVCs) are statically defined across a series of switches and never change. Switched visual circuits (SVCs), require an additional layer 3 route discovery and circuit setup extension to the protocol. Frame relay PVCs are uniquely addressed on switches using data link connection identifiers, (DLCIs). A DLCI is the equivalent of a MAC address on Ethernet or token ring, except that it is only significant to a circuit on a switch, and does not need to be unique across the entire network.

Diagram 1, *Frame relay Network Components*, illustrates these frame relay building blocks.

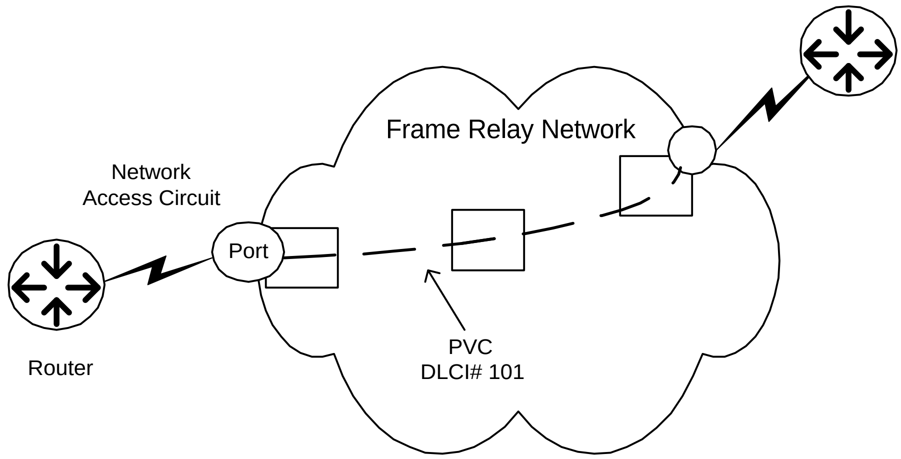

Frame Relay Network Components
Diagram 1

Frame Relay PVC Topology

This section will focus on permanent virtual circuit (PVC) deployments on frame relay networks.

PVC Meshing on a Network

Engineers have many options when deciding on a PVC topology or mesh. They can provision a single PVC to a central circuit, (partially meshed, hub and spoke), or multiple PVC to every other circuit or router on the internetwork (fully meshed), or a subset thereof. Once the overlay mesh is determined, PVC CIRs need to be sized accordingly to meet traffic needs. In order to make intelligent decisions for both meshing and CIR sizing, it is absolutely necessary to understand the application characteristics and WAN traffic patterns on the network. Two typical topologies with their applications and router modeling are described below.

Partially Meshed Topology

Partially meshed topology is the most common and flexible frame relay PVC architecture, and is often implemented as a hub-and-spoke. The hub-and-spoke network is the minimalist approach to frame relay networking, as PVCs from "spoke" circuits are provisioned solely to the central "hub" router circuit(s), and not to every other spoke. Traffic from spoke to spoke traverses the hub circuit(s) and utilizes the hub router as a central switching point. This minimizes the number of PVCs needed, and is a method of achieving maximum economy of scale from a frame relay network. Hub-and-spoke topologies are the best choice for hierarchical application environments where traffic patterns typically flow from remote sites (branch offices) to central (headquarters/data center) facilities. An example is an SNA environment where cluster controllers at branch offices need connectivity to an IBM host in the data center. Diagram 2 illustrates a partially meshed, hub-and-spoke PVC topology.

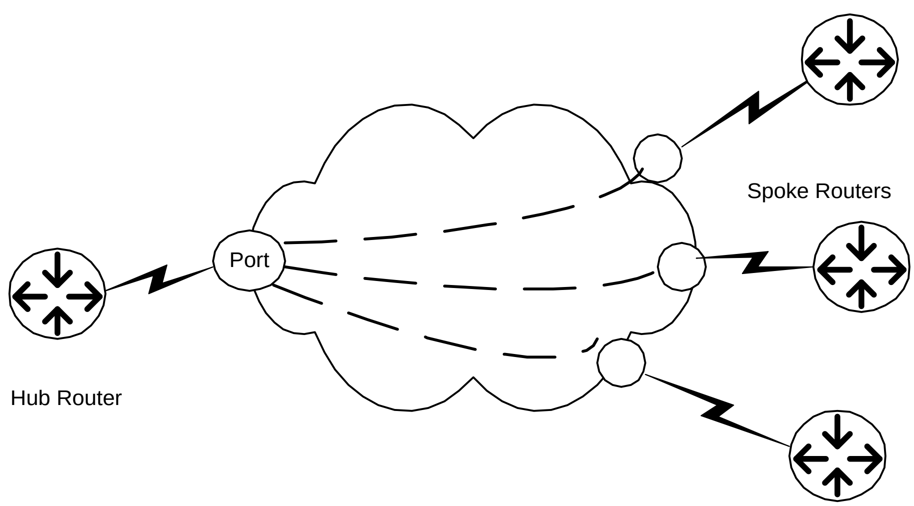

Partially-Meshed, Hub-and-Spoke PVC Topology
Diagram 2

Router Implementations with Partially Meshed PVC Architectures

The hub router in Diagram 2 can either model the three frame relay PVCs as a collection of individual point-to-point circuits (addressing each out of discrete subnets), or as part of a multipoint non-broadcast segment (addressing out of a common subnet). Each method has specific applications depending on upper layer protocol requirements and is implemented through the creation of logical interfaces, which on Cisco routers are known as *subinterfaces*. The Bay Networks router equivalent to the subinterface is the *direct mode PVC*, while 3COM uses *virtual ports*. This paper refers to the concept of virtual frame relay interfaces as *subinterfaces* for consistency.

Subinterfaces
Cisco routers can be configured to "terminate" PVCs on individual subinterfaces that allow a single serial port with multiple PVCs to be treated as multiple logical interfaces under the main interface. Subinterfaces are further defined as either *multipoint* or *point-to-point* on Cisco routers. Diagram 3 illustrates subinterfaces S0.1, S0.2 and S0.3 defined under a physical serial port, S0.

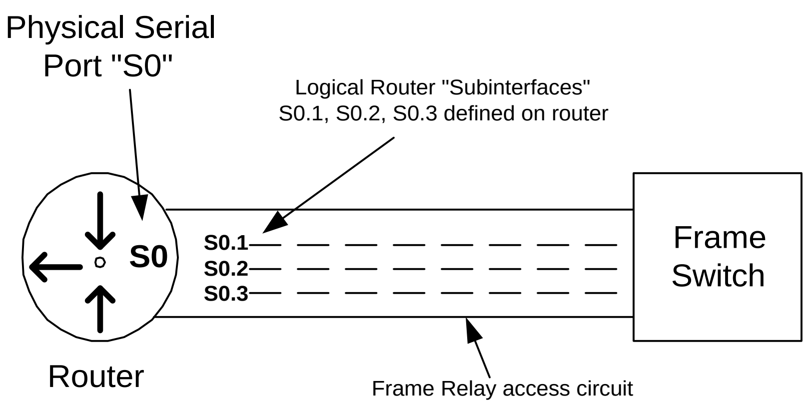

Subinterfaces
Diagram 3

Point-to-Point Subinterfaces
This subinterface type is logically modeled as point-to-point physical interface, like a serial port. With point-to-point, one subinterface per PVC is created and network addressing is allocated out of a discrete, unique subnet. The point-to-point model works well with protocols subject to the limitations of distance vector routing protocols (like IP/IPX RIP) which employ split-horizon to avoid routing loops and speed convergence. Split-horizon is a distance-vector routing protocol characteristic that prevents routing updates for specific networks from being advertised out interfaces from which they were received. This is a problem with frame relay networks where many PVCs terminate under a common interface. Since point-to-point subinterfaces model each PVC as a separate (sub) interface, split horizon is not an issue, and routing updates learned from spokes are copied and propagated across all PVCs by the hub router. In addition, the hub router IP address is advertised as the next-hop address to each spoke router for every destination network. This is important in partially meshed frame relay networks because spoke routers cannot use ARP resolution to determine other spoke routers Layer 2 addresses (DLCI #) if a direct PVC is not provisioned between the two. The point-to-point model also allows optimum control over traffic engineering by allowing manual interface costing, i.e., access control through filtering and frame relay traffic shaping, all on a PVC per PVC basis.

One disadvantage of point-to-point subinterfaces is their excessive consumption of router resources. Since routers treat subinterfaces logically as separate interfaces, they allocate interface buffers to manage the switching of packets to and from each interface. This is sometimes a concern where large numbers of PVCs terminate on hub routers with limited memory.

Diagram 4 illustrates how to logically address point-to-point subinterfaces.

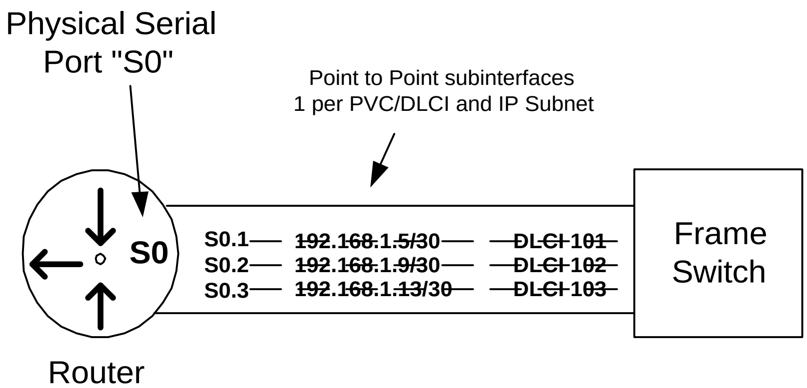

Point to Point Subinterfaces
Diagram 4

Multipoint Subinterfaces
This subinterface type is modeled, logically as a physical interface on a multi-access network, such as a LAN. Multipoint subinterfaces allow all routers on a meshed cloud to be addressed out of a common IP subnet. The Bay Networks equivalent is *group mode* while 3COM handles them as normal frame relay *ports*. Since only a single IP subnet is needed per mesh, rather than one per PVC, this model conserves IP address space. Hub routers using multipoint subinterfaces also have a reduced memory burden since they inherently require fewer numbers of subinterfaces and reserved interface buffers. The multipoint model works well with many implementations of link-state routing protocols, but not distance-vector routing protocols such as RIP due to the behavior of split-horizon as described above. Diagram 5 illustrates the multipoint subinterface concept.

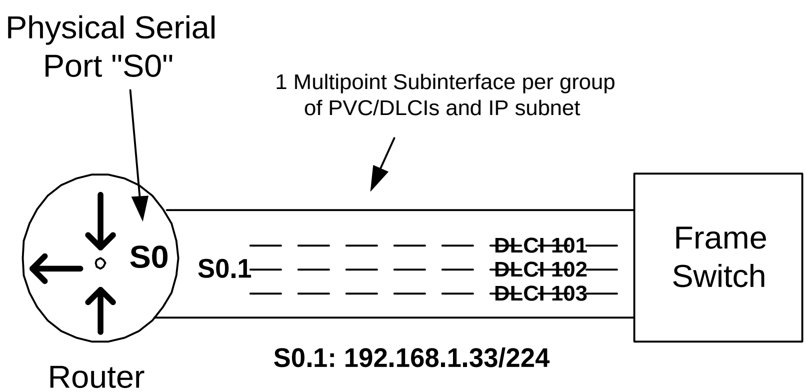

Multipoint Subinterfaces
Diagram 5

Fully Meshed Topologies

In peer-to-peer networking environments where equally dispersed or random traffic patterns are normal, it is often necessary for every router to have PVCs provisioned from its local circuit to every other router's circuit in the WAN. This is known as a *fully meshed* topology. An example of this type of traffic pattern is in a decentralized client/server environment where traffic patterns are unpredictable due to business units being geographically dispersed. Routers comprising a WAN core are also often fully meshed to provide the most direct paths with the least amount of latency out to the distribution layer. This is important to performance with delay-sensitive interactive data traffic such as LAT or SNA. An abundance of router hops and excessive PVC latency can degrade performance and cause timeouts with these protocols. Diagram 6 illustrates fully meshed PVC topologies.

Due to the exceptional number of PVCs required to create them, fully meshed topologies become unwieldy on large networks. The exact number of PVCs needed for a full mesh can be determined by the formula: N x (N-1)/2, where N represents number of circuits. For example, a network with 15 routers, each with one access circuit into a frame relay network will require a total of 105 PVCs for a full mesh. (15 x 14 /2 =105). That same 15-node network would only require 14 PVCs if a hub-and-spoke (partial) mesh were used instead.

Fully-Meshed PVC Topology
Diagram 6

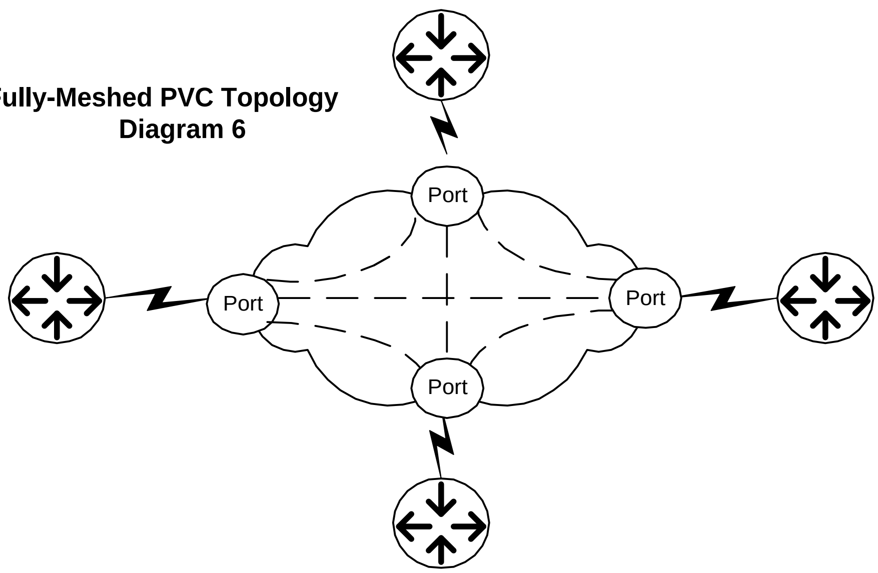

Router Implementations of Fully Meshed PVC Architectures

Both point-to-point and multipoint subinterfaces can be employed on fully meshed frame relay architectures. It is common practice to use multipoint subinterfaces, however, as fully meshed topologies most closely resemble LANs, where every end station can communicate directly with every other at the data link layer. All routers in a multipoint "cloud" are assigned network addresses out of a common subnet, again like in a LAN. Distance vector routing protocol characteristics, such as split horizon, are not problematic with multipoint and fully meshed networks since all routers receive routing updates directly over neighboring PVCs. However, multipoint is not always the model to choose on full meshes. It is important to understand that like a LAN interface, multipoint interfaces assume all nodes attached have access to all bandwidth (frame relay port speed) and do not make provisions for varying PVC CIRs. This becomes a challenge when manual route costing is required on an interface basis to force traffic down a particular path.

OSPF Background

OSPF is an IP routing protocol initially developed by the Internet Engineering Task Force (IETF) in 1987. Extensions to the original specification (OSPFv2) continue to be developed, as this protocol is continually expanded to meet the changing scope of IP networks. OSPF is considered to be an Interior Gateway Routing Protocol (IGP), meaning a routing protocol normally implemented on a network under the control of a single administrative authority. Other IGP examples are RIP versions 1 and 2, Cisco Systems' IGRP and EIGRP, as well as IS-IS. OSPF has fast become the TCP/IP IGP of choice on most mid- to large-sized networks, due to its standards-based deployment, scalability, and fast convergence around link failures.

OSPF is a link state routing protocol, where routers within a common "area" all maintain identical copies of the network map, or topology database. This is accomplished by exchanging the status of networks and "links," which refer to physical router interfaces that connect them. Re-routing around link failures is much quicker with link state routing compared to distance vector protocols like RIP. In a RIP environment, routers have to wait for the route to be unavailable, and go into holddown state before it is flushed, a process that can take a few minutes. In OSPF, all routers constantly maintain an updated topology database and the shortest path to a particular destination can be computed on a router-by-router basis by running a shortest path first routine known as the Dijkstra Algorithm, resulting in quick additions to the routing table.

Given all of its advantages, OSPF is often misunderstood, and presents certain challenges to implement correctly. Poor deployments of OSPF often result in worse network performance than running distance vector routing protocols like RIP. Deploying OSPF is especially challenging on networks that are already installed and addressed, or migrating from another routing protocol. Special challenges and constraints also apply when deploying OSPF on frame relay networks, which is the focus of the following sections.

OSPF Basics

OSPF Neighbors

When routers running OSPF initialize, they attempt to locate neighboring routers to exchange link state advertisements (LSAs) from which routing tables are constructed. Routers form adjacencies with neighboring routers before exchanging this routing information. Details such as subnet address, OSPF area number, network type, and authentication passwords are all checked before an adjacency is formed between neighbors. On broadcast or point-to-point segments neighbor discovery is done dynamically through the ASPF multicast, 224.0.0.5, using the OSPF "hello" protocol. On NBMA networks neighbors must be configured manually before the Hello protocol will initialize in unicast fashion and begin the adjacency forming process.

OSPF Network Types

OSPF has defined standards for communicating across a diverse set of network media.

Broadcast

This OSPF network type typically runs on multi-access broadcast interfaces such as Ethernet, Token Ring, or FDDI. With OSPF broadcast, a designated router (DR) and backup designated router (BDR) are elected dynamically on a broadcast segment, to which all other routers form adjacencies with, to exchange link state information. Election criteria include router ID, loopback interface presence, and router interface priority values, all of which can be manually configured to influence the selection process. It is the responsibility of the DR and BDR to collect link state information from all routers on the broadcast segment, and then compile and distribute the resulting area map back to each. This precludes all routers on a common segment from exchanging link state information with every other router on a segment and reduces the amount of traffic on a broadcast segment. It is important to understand that there is a singular DR/BDR on every broadcast segment in an OSPF network, not just one per area.

The broadcast network type can also run on NBMA media such as frame relay if routers are configured to copy broadcasts across all provisioned PVCs. Frame relay networks should be fully meshed to implement this network type, otherwise DR/BDR confusion may occur. In hub-and-spoke partially meshed architectures, the spokes should be configured so that they never initialize as DR/BDR during elections, since they cannot form adjacencies with other spoke routers. This is done by setting the OSPF priority on the frame relay subinterface to zero.

Point-to-point (Pt-to-Pt)

This OSPF network type is typically implemented across dedicated WAN circuits such as T-1 links or on frame relay point-to-point subinterfaces. No designated router is elected on point-to-point networks, since only two routers exist on a segment. They exchange link state information and routes as peers across the common subnet. This is the default network type for point-to-point frame relay subinterfaces.

Non-Broadcast Multi-Access (NBMA)

This network type was developed for OSPF to run on media such as X.25, frame relay, and ATM, where the network cannot dynamically forward broadcast packets to all other routers in a virtual network. Other than manual configuration of OSPF neighbors, router behavior configured with this network type is identical to that of broadcast. (I.e., Hello protocol elects DR/BDR and adjacency with all non-DR/BDR routers are formed.) Again, it is important to ensure that a hub router is elected to be the DR on a hub-and-spoke partially meshed frame relay network to ensure that adjacencies can be formed with every spoke.

Point-to-Multipoint (Pt-to-Mp)

This network type was developed for OSPF to run on NBMA networks such as Frame relay and ATM. Routers are addressed out of a common IP subnet on WAN interfaces but full meshing is not required, as DR/BDR election is not done on point-to-multipoint segments. This network type is well suited for frame relay hub-and-spoke networks where conservation of IP addresses or minimizing resource impact of logical interfaces on hub routers is an issue. Any-to-any spoke connectivity in a partially meshed PVC environment is possible since the hub router will advertise itself as the next hop forwarding address to all spokes for all routes.

OSPF Areas

OSPF was constructed to be hierarchical in nature, meaning domains are partitioned into separate manageable groups known as areas. Medium to large networks perform best when routers are grouped into multiple areas, each of which are laid out after careful consideration to IP addressing and application traffic patterns are taken. Guidelines for the number of routers per area and other related issues will be discussed in later sections.

Backbone Areas

When more than one area exists in a network, a “backbone” area referred to as Area Zero must be constructed, which serves as the transit, or relay area between all other areas. Placement of this backbone area within the internetwork is a crucial design consideration that will be discussed in later sections.

Non-Zero Areas

Any area other than the backbone area must be defined as a non-zero area, that is, represented with an integer (Area 10) or dotted decimal (Area 10.10.10.10) notation. Every router within a non-zero area will maintain an identical copy of the area topology database, and recompute the Dijkstra Algorithm during times of link failure. Traffic between networks in the same non-zero area will always traverse routers within the same area if a valid path exists.

Area Border Router

Area border routers (ABRs) touch the backbone “area” and any other non-zero area. These routers maintain multiple link state topology databases, specifically one per area that they connect to. It is the responsibility of an ABR to forward summary information from one area to another upon startup and during periods of convergence. Route summarization in OSPF networks is accomplished by utilizing ABRs as route aggregation points and announcing summaries into the backbone area.

OSPF Internal/ Intra-Area and Inter-Area Routes

All routes sourced from within an OSPF domain are considered to be *internal* routes. IP routes known to a router that originated in the same Area are known as *intra-area* OSPF routes, where those that originated in a different area are known as *inter-area* OSPF routes.

OSPF External Routes/Autonomous System Boundary Routers

External routes originate outside of the local OSPF domain. Autonomous system boundary routers (ASBRs) are routers that partake in a routing process outside of and in addition to the local OSPF process, and exchange between the two. For example, a router that runs RIP on certain interfaces and OSPF on others would be an ASBR if it exchanged routing information between the two domains. Routes learned from the RIP domain would be defined as OSPF external routes when propagated into the OSPF domain by the ASBR. Placement of ASBRs within a network should be carefully considered, since it will impact the type of OSPF area that can be deployed in the areas where they reside. This will be discussed below in the Stub Area section.

Stub Areas

Stub areas require only a subset of the complete routing table. For example, networks with only limited paths out to the rest of the internetwork (i.e., spoke routers in a hub-and-spoke frame relay topology) can be defined as stub areas. ABRs that connect to stub areas will flood specific routes from only networks within that same area. A default route (0.0.0.0) will be generated in lieu of all external (and sometimes inter-area internal) routes. This can greatly reduce the amount of memory consumption on stub routers by limiting the size of routing tables and number of link state databases. ASBRs cannot exist in normal stub areas as they inherently produce external routes. Diagram 7 illustrates OSPF components.

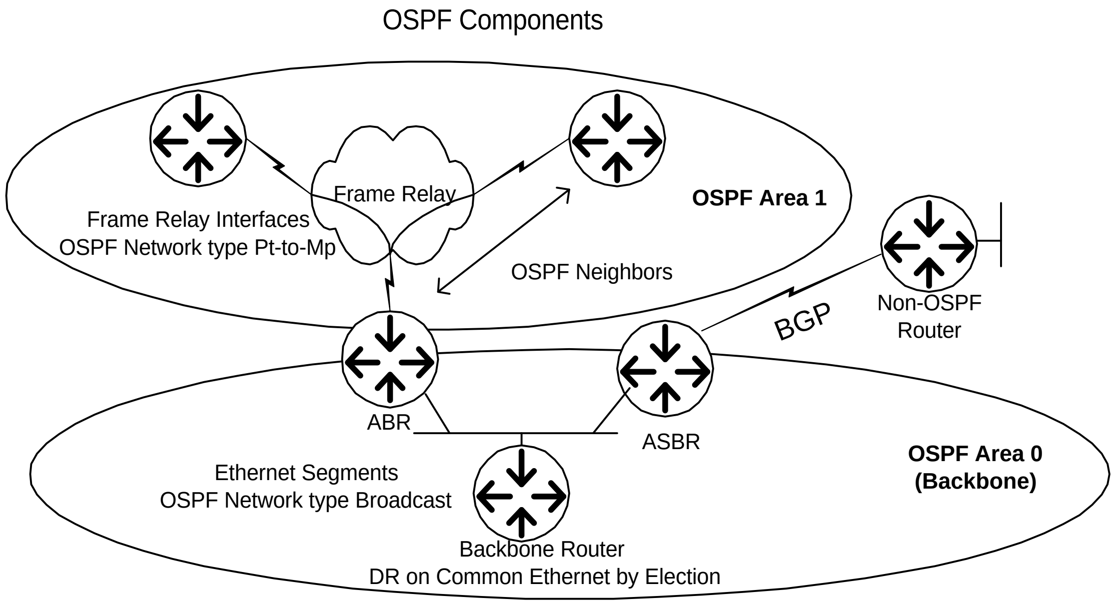

DIAGRAM 7 OSPF COMPONENTS

Design Methodology

The cleanest way to implement OSPF on frame relay networks is to design and build everything from scratch. Design decisions concerning network addressing and PVC provisioning can be made soundly and confidently based purely on business needs and application requirements. Assuming a reasonable budget, it is straightforward to design and implement a small- to medium-size IP network with OSPF on frame relay.

Adding OSPF to an existing frame relay network presents additional challenges depending on the existing deployment of IP addressing, PVC architecture, the installed router base, and other non-IP protocols deployed. Regardless of the situation, it is necessary to use a consistent approach when designing and implementing OSPF on frame relay networks.

Requirements Gathering

Flushing out the requirements of a new network is often an excruciating, lengthy task. This is far from being wasted time, however, since being forced to redesign a newly implemented network due to a major oversight is both time-consuming and embarrassing. Interview as many of the network managers, engineers, operations staff and vendors as possible to identify key issues that will impact a new network.

Business Factors

How critical is the application data to the company? Is 100% up time during production hours necessary as in a stock trader or manufacturing network? If this is the case, redundancy and dial-backup will most likely be an integral part of the design. What is the company business model? Do branch offices report to districts and then to headquarters? What are the budget constraints? Will a full-scale equipment and circuit rollout occur at implementation, or will it be a phased approach due to budgeting constraints?

Network Application Characteristics

What is the nature of the application traffic? Are the applications based on a client-server model, batch processes on a mainframe, or terminal emulation on a VAX or AS-400 minicomputer? What are the expected traffic patterns of the applications? Are servers located at distributed district offices or at a centrally located data center? Do all spoke router locations need to communicate directly to every other location or just to the hub networks? Are any non-IP protocols required to be routed or bridged natively across this network?

Vendor Decisions: Will this network be outsourced for implementation and management or will this be done internally? Will private or public frame relay infrastructure be utilized? What hardware will be or is currently used on this network?

Network Scale: How many remote sites and networks will connect to the network? What is the breakdown of sites within a particular business unit or district? How many hosts are there per network? Are there third-party networks requiring connectivity, or will this network be subject to integration into a larger internetwork?

Current IP Addressing: For existing networks, how is IP addressing deployed? If deployed in a flat, non-hierarchical fashion, is it possible to readdress? Can secondary IP addresses be used to transition from a flat to hierarchical structure to support route aggregation? Is the network small enough to leave flat if OSPF is only being implemented for VLSM support? What is the routing protocol deployed on the campus network and will redistribution be necessary?

Network Design

A comprehensive design should be a response to the requirements that were initially gathered. The design should be both a tool for installation engineers to use during the implementation, as well as a source of documentation for ongoing maintenance and management of the network. It is often useful to include the original requirements in the text of a design to help determine when changes to scope may warrant a redesign. A good design includes a high-level description of network applications, schematics of physical and logical connectivity, spreadsheets documenting circuit information and IP addressing, as well as equipment specifications.

Hardware Selection

Specify router hardware for hub-and-spoke sites. Consider the use of routers with integrated DSU/CSU interfaces to minimize cabling and eliminate potential points of failure. The goal is to avoid WAN circuit flapping, which is a major source of area instability in OSPF networks. Specifying memory for routers is sometimes a difficult task. In general, load hub routers with extra memory to facilitate ABR functionality, as well as copying broadcasts across PVCs and handling OSPF neighbor exchanges. Include redundant hub routers in design in case of a mission critical network requirement. Include ISDN or asynchronous terminal server(s) for dial-backup if needed.

Frame Relay Mesh and Circuit Sizing

Size circuit port sizes and CIRs according to estimated traffic requirements. If possible, try not to oversubscribe the frame relay MIR more than 200 – 250% in regards to remote aggregate vs. hub port sizing. Avoid oversubscribing CIR whenever possible for consistent behavior during peak loads. Networks with applications having a low tolerance for delay should be provisioned more conservatively. Consider the use of partially meshed PVC structures for simple hub-and-spoke networks unless peer-to-peer application traffic flows are commonly implemented. If the network is outsourced to a service provider for management, consider installation of a “firewall” router between the “managed” frame relay hub router and

Router Frame Relay PVC Modeling

Point-to-Point

When in doubt, employ the point-to-point PVC models on routers when implementing hub-and-spoke architectures. (Define frame relay interfaces as point-to-point, direct mode PVC, or virtual port, depending on router vendor.) This model consumes the most IP addressing and router resources to implement but gives the most granularity for traffic engineering and is the easiest to troubleshoot. This model is often the only choice for multi-protocol networks, which employ routing protocols subject to the constraints of split-horizon (i.e., IPX, AppleTalk, and DECnet).

Multipoint and Group Mode

One practical advantage of this model is that it conserves IP address space by reducing the number of discrete subnets needed in a large network. This may be an issue where a diminishing pool of registered IP addresses is available. (Consider IP unnumbered or “private” WAN addressing if this is the only compelling reason to implement the multipoint model.) Another advantage is that the multipoint model eases the burden on hub routers by reducing the number of required interface buffers, routing table/link state database size, and broadcast propagation. This model works well in networks where required connectivity is limited to simple spoke to hub router and maximizing hub router performance is of concern. A practical example is a network using TCP/IP to carry DLSw traffic between routers from remote SDLC or LLC2 devices to a central SNA host. Since connectivity is limited from hub to spoke and DLSw buffering consumes significant router resources, the multipoint model is well suited.

Non-Zero Area
Place frame relay interfaces into specific non-zero areas based on geographical or business unit factors. Whenever possible, put routers connecting networks with frequent traffic flows into the same OSPF areas. Cisco’s recommendation of no more than 30 to 40 routers per (non-zero) area suites the hub-and-spoke model well and can be scaled to at least 50 to 60 routers depending on router memory, stub area deployment and application traffic resource demands.

Extended Hub-and-spoke, Two-Tier Networks
For a national company with a regional branch/district reporting structure, it is often advantageous to extend the OSPF backbone across the WAN, out to the district routers. This allows the OSPF hierarchy to follow the organizational structure of the business more closely, where the districts are all peers and branches report to districts. Each district router becomes an ABR in this instance, and branch routers are treated as spokes.

Backbone Area
Form the OSPF backbone using the common WAN Circuits/PVCs between district routers. Consider redundant circuits and fully meshed PVCs to form a resilient core and maximize OSPF Area 0 stability.

Non-Zero Area
Place branch routers under a district office into common non-zero OSPF areas. This architecture facilitates efficient intra-area routing within districts, since routing will never traverse Area 0.

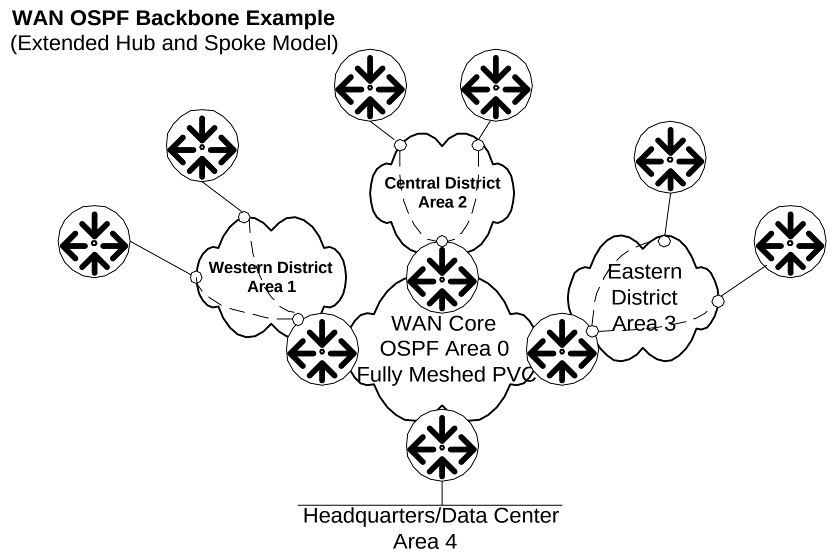

Flat OSPF Hierarchy
In extremely small networks (30 routers or less), or network trials, it is often advantageous to include all routers in the same area to construct an OSPF network with no hierarchy. It is a good idea to create this area as a non-zero area to allow for expansion and/or incorporation into a corporate OSPF network at a later time. This would prevent area renumbering to remove the "backbone" off of the WAN routers.

OSPF Network Types

WAN
Frame relay interfaces on Cisco routers can be configured for any of the four supported network types, depending how the PVC meshing is modeled. Point-to-point/direct mode PVC modeling will always default the OSPF network type to point-to-point. If implementing multipoint interfaces with partial meshes, it is best to implement OSPF network type point-to-multipoint. This will preclude a DR/BDR election process and facilitate any-to-any connectivity. Fully meshed PVC structures which model a broadcast LAN can be configured for the broadcast or NBMA network types. Since these models employ an OSPF DR/BDR, it is important that the full mesh always be present so that every router forms adjacencies with these routers to provide consistent information into the link state databases. In the above example of a WAN core configured as a backbone area, it would be advisable to provide redundant access circuits from each router with a full mesh to every other router to ensure maximum backbone integrity during PVC failures.

LAN
Ethernet, Token Ring, and FDDI interfaces on routers will default to the broadcast network type.

Designated Router
Routers with OSPF network types set as broadcast or NBMA will participate in designated router (DR) and backup designated router (BDR) election to control distribution of LSA information. It is a good idea to prioritize interfaces in such a manner that routers connected to multiple LAN segments do not get elected as the DR or BDR on all of the segments. Instead, distribute this functionality across multiple routers in a network with many broadcast segments to balance the load.

IP Addressing
The power of OSPF can only be fully leveraged if IP addresses are assigned in contiguous block to routers and hosts within common areas. In an area, routers send network advertisements with full-length subnet masks, exactly as they receive them. It is the ABRs' responsibility to collect and summarize contiguous blocks of networks into the backbone area by using advertisements with shorter subnet masks. This greatly reduces the size of routing tables within large networks, which subsequently reduces load on router resources and speeds network convergence when links fail. OSPF summarization is not enabled by default, and requires manual configuration on an ABR.

Route Summarization on New Networks
New networks can be easily architected to employ route summarization to the fullest degree. A common mistake is to assign all WAN links out of one common subnet and deploy them on routers across multiple areas. This limits your ability to fully summarize the WAN address space, and may require you to announce each link subnet explicitly. Instead, assign contiguous blocks to an area and then develop the subnetting plan for WAN and LAN networks from that same space. This facilitates the cleanest summarization in an OSPF network.

The following graph illustrates the benefits of route summarization.

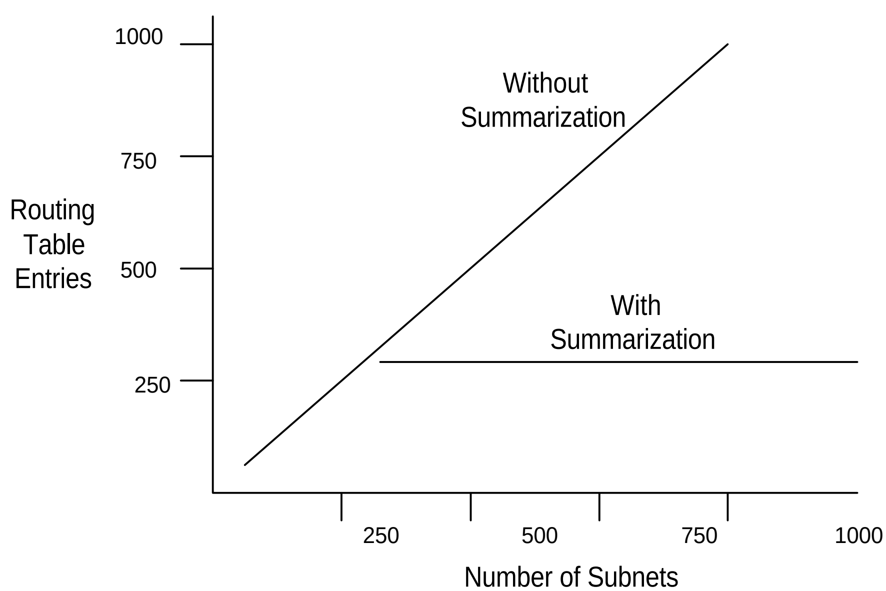

Route Summarization on Existing Networks
Even on networks with flat topologies and scattered IP addressing, it is often possible to summarize small groups of routes to minimize the size of routing tables. This is done by identifying contiguous blocks of routes that can be summarized on even bit boundaries and then leaking the specific routes that fall outside of a summarization block into the backbone.

Consider the following example:

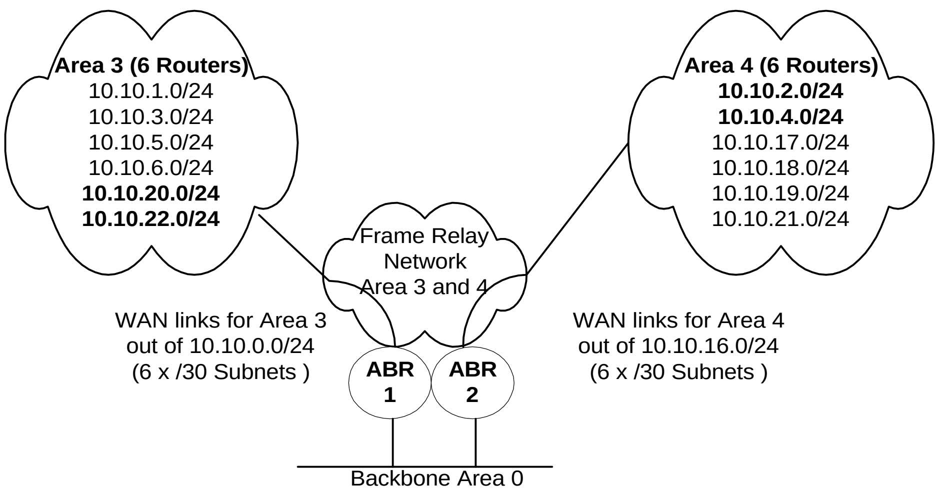

Without route summarization, each ABR will pass **6 -** /24 network advertisements and **6 -** /30 network advertisements towards the backbone, for a total of **24 network advertisements** onto the backbone from both ABRs.

It is possible to reduce this to **6 network advertisements** onto the backbone by summarizing as follows:

ABR 1
Router ospf 100
Network 10.10.0.0 0.0.0.255 area 3 (This places the WAN links on ABR1 into OSPF Area 3)
Area 3 range 10.10.0.0 255.255.248.0 (Turns on route summarization for Area 3)
!
ABR 2
Router ospf 100
Network 10.10.16.0 0.0.0.255 area 4 (This places the WAN links on ABR2 into OSPF Area 4)
Area 4 range 10.10.16.0 255.255.248.0 (Turns on route summarization for Area 4)
!

ABR1 will advertise a summary of 10.10.0.0/21, and specific routes of 10.10.20.0/24 and 10.10.22.0/24
ABR2 will advertise a summary of 10.10.16.0/21 and specific routes of 10.10.2.0/24 and 10.10.4.0/24

By summarizing where possible, and leaking specific routes where necessary, OSPF networks can be optimized even in situations where less than perfect IP address allocation has been deployed.

Proof of Concept
Whenever possible, it is beneficial to implement a portion of the network as part of an application pilot or proof of concept phase. For example, the hub router and one or two spokes can be implemented to measure client/server response time of a SNA application. It is much easier to identify and effect application tuning parameters on a lightly loaded network rather than once it is fully rolled-out and considered production. SNMP-based trend analysis tools measure circuit utilization consumed by applications and adjust CIRs for future sites when needed. Network simulation software is the next best thing to an actual pilot, and is useful in making many "what-if?" design decisions.

Implementation Plan
A network design is worthless by itself if poorly implemented. Attention to detail during the installation phase of an OSPF network is crucial to the success of any decent sized or aggressive router rollout. Carefully constructed installation and test plans with pre-built router configuration files maximize the chances of a smooth implementation. Design documentation should contain all of the addressing and customization details that router engineers need in a readily available format. One possible format is illustrated below:

Spoke Router	DLCI # To S.F. Hub	WAN IP (/30 prefix)	LAN IP (/24 prefix)	OSPF Area	WAN Cost	LAN Cost
Seattle	200	10.16.254.1	10.1.1.0	2	1500	default
Portland	200	10.16.254.5	10.1.4.0	2	1500	default

OSPF areas should be sized in anticipation of network growth (within existing areas), and modular configurations constructed that can be seamlessly expanded during phased implementations. For example, it is often a good idea to place all of the spoke routers within an OSPF area onto a common hub router circuit when provisioning PVCs. This facilitates a smooth expansion with minimal configuration changes (none on the spokes), if redundant or more hub routers are installed, since capacity is required during a phased implementation.

Summary

In conclusion, it is obvious that designing frame relay networks using OSPF is not a trivial task and requires knowledge of the protocols, applications, and user requirements. A summary of major points follows:

- Frame relay is a good choice for WANs requiring speeds of 56KB up to NxT-1 and DS-3
- OSPF is a good fit for TCP/IP based frame relay networks requiring VLSM support and route summarization
- Thorough requirements gathering is crucial to designing sound, scalable networks. Determine application traffic flows before deciding on frame relay mesh and OSPF architecture.
- Develop detailed design documentation to facilitate smooth installs and ongoing operations of networks.
- Architect frame relay mesh to model application flows and size circuits and CIRs accordingly.
  - Oversubsription of remotes to hub port speeds should not exceed 200 – 250%
  - Avoid oversubsription of CIR when application traffic loads are uncertain.
- Implement point-to-point modeling of PVCs on routers when granular route costing is necessary or multi-protocol traffic is required.
- Implement point-to-multipoint modeling of PVCs on routers when simple hub-and-spoke connectivity is required and address conservation or hub router memory is a concern.
- Architect OSPF area layouts, once application traffic patterns are understood and frame relay mesh is determined.
- Placement of the backbone area is crucial to a network of optimal design.
  - For simple hub-and-spoke networks place LAN interface(s) of hub router into the backbone Area 0. Do not extend the backbone across the WAN.
  - For regional extended hub-and-spoke networks, create a WAN backbone Area 0 with redundant circuits or PVCs.
- Organize non-zero areas in manageable groups mirroring traffic patterns.
  - Limit OSPF areas to 30-40 routers when possible.
  - Limit ABRs to 2-3 areas.
- Assign IP addressing to OSPF areas in contiguous blocks to allow summarization into the backbone.
- Conduct application pilots or of partial OSPF network to revise design before full-scale rollout.
- Develop detailed implementation plans to provide consistent, accurate router configurations.

References

OSPF Design Guide, http://www.cisco.com/warp/customer/104/1.html#I00, Cisco Systems Inc.

Routing in The Internet, Christian Huitema

*OSPF***,** *Anatomy of an Internet Routing Protocol,* John T. Moy

*Cisco IOS Software Command Summary/Configuration Guide***,** Cisco Systems Inc.

OSPF Network Design Solutions, Thomas M. Thomas II